

WEST Search History

DATE: Monday, April 16, 2007

Hide?	<u>Set</u> <u>Name</u>	<u>Query</u>	<u>Hit</u> <u>Count</u>
	<i>DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=NO; OP=OR</i>		
<input type="checkbox"/>	L24	(l20 or l21 or l22 or l23) and l18	2
<input type="checkbox"/>	L23	PICKEL-JAMES-WILLIS.in.	8
<input type="checkbox"/>	L22	JOSTEN-JEFFREY-WILLIAM.in.	11
<input type="checkbox"/>	L21	HOWELL-BRIAN-K.in.	5
<input type="checkbox"/>	L20	COTNER-CURT-LEE.in.	22
<input type="checkbox"/>	L19	L18 and (l5 or l6)	4
<input type="checkbox"/>	L18	(indoubt near entr\$)	4
<input type="checkbox"/>	L17	indoubt	53
<input type="checkbox"/>	L16	((transaction adj1 manager) with (database\$ or (data adj1 base\$)) with recording with entries with (rollback or commit) with (shared adj1 memory))	0
<input type="checkbox"/>	L15	(l10 or l11 or l12 or L13) and l6	285
<input type="checkbox"/>	L14	(l10 or l11 or l12 or L13) and l5	37
<input type="checkbox"/>	L13	707/202-204.ccls.	4019
<input type="checkbox"/>	L12	707/200.ccls.	2925
<input type="checkbox"/>	L11	707/10.ccls.	6965
<input type="checkbox"/>	L10	707/1.ccls.	5712
<input type="checkbox"/>	L9	((shared adj1 memory) with (rollback or commit) with (database\$ or (data adj1 base\$)) with recording with failure with (transaction adj1 manager))	0
<input type="checkbox"/>	L8	((shared adj1 (database\$ or (data adj1 base\$))) with recording with transaction)	0
<input type="checkbox"/>	L7	((shared adj1 (database\$ or (data adj1 base\$))) with recording with transaction with (rollback or commit))	0
<input type="checkbox"/>	L6	(two near phase near commit)	1178
<input type="checkbox"/>	L5	"2-phase commit"	107
	<i>DB=USPT; PLUR=NO; OP=OR</i>		
<input type="checkbox"/>	L4	L3 and recover\$	3
<input type="checkbox"/>	L3	L2 and (rollback or (roll adj1 back) or roll-back or commit)	3
<input type="checkbox"/>	L2	L1 and (database\$ or (data adj1 base\$))	3
<input type="checkbox"/>	L1	"2-phase commit".ti.	3

END OF SEARCH HISTORY

10/726,700


[Web](#) [Images](#) [Video](#) [News](#) [Maps](#) [more »](#)

[Advanced Scholar Search](#)
[Scholar Preferences](#)
[Scholar Help](#)

The "AND" operator is unnecessary -- we include all search terms by default. [\[details\]](#)

Scholar [All articles](#) [Recent articles](#) Results 1 - 10 of about 1,810 for **database and 2 phase commit and rc**

All Results

[C Mohan](#)
[C Pu](#)
[H Berenson](#)
[G Weikum](#)
[D DeWitt](#)

[Data processing method for file status recovery includes providing a log file of atomic transactions ... - group of 2 »](#)

AH Frey Jr, RC Mosteller, JM Gould, NR Mendelsohn, ... - US Patent 5,201,044, 1993 - Google Patents

... No. 4,819,159 to Shipley et al., a fault-tol- The Lindsay et al **database** employs a distributed trans- ... In effect, this accomplishes a two ... Thesecond **phase** ...

[Cited by 69](#) - [Related Articles](#) - [Web Search](#)

[Distributed optimistic concurrency control with reduced rollback](#)

D Agrawal, AJ Bernstein, P Gupta, S Sengupta - Distributed Computing, 1987 - Springer ... committed and written to the **database**, T1 can ... than they leave the write **phase**, the **transaction** ... overcome this problem MVPV uses two **transaction** counters, ctn ...

[Cited by 34](#) - [Related Articles](#) - [Web Search](#)

[Berkeley DB - group of 7 »](#)

MA Olson, K Bostic, M Seltzer - [usenix.org](#)

... are flushed to the log disk when transactions **commit**. ... DB provides a service known as two-**phase** locking ... and to guarantee ACID properties, **database** systems manage ...

[Cited by 111](#) - [Related Articles](#) - [Cached](#) - [Web Search](#)

[System, method and program for performing two-phase commit with a coordinator that performs no ... - group of 3 »](#)

CL Cotner, RA Crus, BK Howell, JW Pickel, DJ ... - US Patent 5,884,327, 1999 - Google Patents

... retrievals, with the data in the **database** over the ... is an extension of the basic two-**phase** **commit** protocol, and ... 2 shows a **transaction** that aborts 5, followed by ...

[Cited by 35](#) - [Related Articles](#) - [Web Search](#)

[Two-pass multi-version read consistency - group of 2 »](#)

G Hallmark, R Bamford - US Patent 5,452,445, 1995 - Google Patents

... master **database**. The second **phase** of the two-**phase** consistency after a fault and not to a system for provid- **commit** is the actual **commit** step. ...

[Cited by 40](#) - [Related Articles](#) - [Web Search](#)

[A critique of ANSI SQL isolation levels - group of 12 »](#)

H Berenson, P Bernstein, J Gray, J Melton, EO'Neil ... - ACM SIGMOD Record, 1995 - [portal.acm.org](#)

... leak requests, thus de- parting from pure two-**phase** locking ... a set of operations that transform the **database** from one ... a read of x by **transaction 2**. **Transaction 1** ...

[Cited by 149](#) - [Related Articles](#) - [Web Search](#) - [BL Direct](#)

[Method and apparatus for improved transaction recovery - group of 6 »](#)

GC Ngai, H Rizvi, LL Tan - US Patent 5,850,507, 1998 - Google Patents

... transactions that were undergoing a two-**phase** **commit** coord- dinated ... 8". Thus, **transaction**

10/726,700


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide

database and 2 phase commit and rollback and entries

SEARCH

THE ACM DIGITAL LIBRARY


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used

database and 2 phase commit and rollback and entries

Found 87,647 of 199,915

Sort results by

relevance

☒ Save results to a Binder

Try an Advanced Search

Try this search in The ACM Guide

Display results

expanded form

☒ Search Tips

☐ Open results in a new window

Results 1 - 20 of 200

 Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

 Relevance scale ☐ ☐ ☐ ☐ ☐

1 [ARIES: a transaction recovery method supporting fine-granularity locking and partial rollbacks using write-ahead logging](#)

C. Mohan, Don Haderle, Bruce Lindsay, Hamid Pirahesh, Peter Schwarz

 March 1992 **ACM Transactions on Database Systems (TODS)**, Volume 17 Issue 1

Publisher: ACM Press

 Full text available: [pdf\(5.23 MB\)](#)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

DB2TM, IMS, and TandemTM systems. ARIES is applicable not only to database management systems but also to persistent object-oriented languages, recoverable file systems and transaction-based operating systems. ARIES has been implemented, to varying degrees, in IBM's OS/2TM Extended Edition Database Manager, DB2, Workstation Data Save Facility/VM, Starburst and QuickSilver, and in the University of Wisconsin's EXODUS and Gamma d...

Keywords: buffer management, latching, locking, space management, write-ahead logging

2 [B-tree concurrency control and recovery in page-server database systems](#)

Ibrahim Jaluta, Seppo Sippu, Eljas Soisalon-Soininen

 March 2006 **ACM Transactions on Database Systems (TODS)**, Volume 31 Issue 1

Publisher: ACM Press

 Full text available: [pdf\(401.86 KB\)](#)

 Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We develop new algorithms for the management of transactions in a page-shipping client-server database system in which the physical database is organized as a sparse B-tree index. Our starvation-free fine-grained locking protocol combines adaptive callbacks with key-range locking and guarantees repeatable-read-level isolation (i.e., serializability) for transactions containing any number of record insertions, record deletions, and key-range scans. Partial and total rollbacks of client transactio ...

Keywords: ARIES, ARIES/CSA, B-tree, cache consistency, callback locking, client-server database system, data shipping, key-range locking, page server, partial rollback, physiological logging, sparse B-tree, structure modification

10/17/2006, 7:00

3 Garbage collection for a client-server persistent object store



Laurent Amsaleg, Michael J. Franklin, Olivier Gruber

August 1999 **ACM Transactions on Computer Systems (TOCS)**, Volume 17 Issue 3

Publisher: ACM Press

Full text available: pdf(267.18 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

We describe an efficient server-based algorithm for garbage collecting persistent object stores in a client-server environment. The algorithm is incremental and runs concurrently with client transactions. Unlike previous algorithms, it does not hold any transactional locks on data and does not require callbacks to clients. It is fault-tolerant, but performs very little logging. The algorithm has been designed to be integrated into existing systems, and therefore it works with standard i ...

Keywords: client-server system, logging, persistent object-store, recovery

4 The relational model for database management: version 2

E. F. Codd

January 1990 Book

Publisher: Addison-Wesley Longman Publishing Co., Inc.

Full text available: pdf(28.61 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

From the Preface (See Front Matter for full Preface)

An important adjunct to precision is a sound theoretical foundation. The relational model is solidly based on two parts of mathematics: first-order predicate logic and the theory of relations. This book, however, does not dwell on the theoretical foundations, but rather on all the features of the relational model that I now perceive as important for database users, and therefore for DBMS vendors. My perceptions result from 20 y...

5 DLFM: a transactional resource manager



Hui-I Hsiao, Inderpal Narang

May 2000 **ACM SIGMOD Record , Proceedings of the 2000 ACM SIGMOD international conference on Management of data SIGMOD '00**, Volume 29 Issue 2

Publisher: ACM Press

Full text available: pdf(124.99 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The DataLinks technology developed at IBM Almaden Research Center and now available in DB2 UDB 5.2 introduces a new data type called DATALINK for a database to reference and manage files stored external to the database. An external file is put under a database control by "linking" the file to the database. Control to a file can also be removed by "unlinking" it. The technology provides transactional semantics with respect to linking or unlinking the file when DATALINK ...

6 High level specification of concurrency control in distributed database systems

L. Chiu, M. T. Liu

April 1988 **Proceedings of the 10th international conference on Software engineering ICSE '88**

Publisher: IEEE Computer Society Press

Full text available: pdf(1.00 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Concurrency control is one of the major issues in database systems; therefore, many concurrency control algorithms based on different strategies have been proposed. Unfortunately there is still lack of a general model for describing these algorithms. Hence,

algorithms cannot be uniformly presented, which makes it hard to understand them and to prove their correctness. This paper proposes a high level specification, based on an object-oriented model, of concurrency control algorithms. Concur ...

7 Dynamic voting algorithms for maintaining the consistency of a replicated database



S. Jajodia, David Mutchler

June 1990 **ACM Transactions on Database Systems (TODS)**, Volume 15 Issue 2

Publisher: ACM Press

Full text available: [pdf\(4.07 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

There are several replica control algorithms for managing replicated files in the face of network partitioning due to site or communication link failures. Pessimistic algorithms ensure consistency at the price of reduced availability; they permit at most one (distinguished) partition to process updates at any given time. The best known pessimistic algorithm, voting, is a "static" algorithm, meaning that all potential distinguished partitions can be listed in adv ...

8 Concurrency control and recovery for balanced B-link trees

Ibrahim Jaluta, Seppo Sippu, Eljas Soisalon-Soininen

April 2005 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 14 Issue 2

Publisher: Springer-Verlag New York, Inc.

Full text available: [pdf\(302.02 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [citations](#)

In this paper we present new concurrent and recoverable B-link-tree algorithms. Unlike previous algorithms, ours maintain the balance of the B-link tree at all times, so that a logarithmic time bound for a search or an update operation is guaranteed under arbitrary sequences of record insertions and deletions. A database transaction can contain any number of operations of the form "fetch the first (or next) matching record", "insert a record", or "delete a reco ...

Keywords: Concurrency control, Recovery, Transaction, Tree-structure modifications

9 A structured approach for the definition of the semantics of active databases



Piero Fraternali, Letizia Tanca

December 1995 **ACM Transactions on Database Systems (TODS)**, Volume 20 Issue 4

Publisher: ACM Press

Full text available: [pdf\(4.15 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Active DBMSs couple database technology with rule-based programming to achieve the capability of reaction to database (and possibly external) stimuli, called events. The reactive capabilities of active databases are useful for a wide spectrum of applications, including security, view materialization, integrity checking and enforcement, or heterogeneous database integration, which makes this technology very promising for the near future. An active database system consists of ...

Keywords: active database systems, database rule processing, events, fixpoint semantics, rules, semantics


10 Model and verification of a data manager based on ARIES



Dean Kuo

December 1996 **ACM Transactions on Database Systems (TODS)**, Volume 21 Issue 4

Publisher: ACM Press

Full text available:  [pdf\(813.93 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

In this article, we model and verify a data manager whose algorithm is based on ARIES. The work uses the I/O automata method as the formal model and the definition of correctness is defined on the interface between the scheduler and the data manager.

Keywords: ARIES, I/O automata, system failures

11 Synchronization and recovery in a client-server storage system

E. Panagos, A. Biliris

August 1997 **The VLDB Journal — The International Journal on Very Large Data**

Bases, Volume 6 Issue 3

Publisher: Springer-Verlag New York, Inc.

Full text available:  [pdf\(205.25 KB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

Client-server object-oriented database management systems differ significantly from traditional centralized systems in terms of their architecture and the applications they target. In this paper, we present the client-server architecture of the EOS storage manager and we describe the concurrency control and recovery mechanisms it employs. EOS offers a semi-optimistic locking scheme based on the multi-granularity two-version two-phase locking protocol. Under this scheme, multiple concurrent reads ...

Keywords: Checkpoint, Client-server architecture, Object management, Concurrency control, Locking, Logging, Recovery, Transaction management

12 A Survey of Techniques for Synchronization and Recovery in Decentralized Computer Systems

Walter H. Kohler

June 1981 **ACM Computing Surveys (CSUR)**, Volume 13 Issue 2

Publisher: ACM Press


Full text available:  [pdf\(3.33 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

13 Partitioned two-phase locking

Meichun Hsu, Arvola Chan

December 1986 **ACM Transactions on Database Systems (TODS)**, Volume 11 Issue 4

Publisher: ACM Press

Full text available:  [pdf\(1.31 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

In a large integrated database, there often exists an "information hierarchy," where both raw data and derived data are stored and used together. Therefore, among update transactions, there will often be some that perform only read accesses from a certain (i.e., the "raw" data) portion of the database and write into another (i.e., the "derived" data) portion. A conventional concurrency control algorithm would have treated such transactions as regular updates ...

14 Concurrency control: methods, performance, and analysis

Alexander Thomasian

March 1998 **ACM Computing Surveys (CSUR)**, Volume 30 Issue 1

Publisher: ACM Press

Full text available:  [pdf\(427.18 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Keywords: Markov chains, adaptive methods, concurrency control, data contention, deadlocks, flow diagrams, load control, optimistic concurrency control, queueing network models, restart-oriented locking methods, serializability, thrashing, two-phase locking, two-phase processing, wait depth limited methods

15 Highly concurrent cache consistency for indices in client-server database systems



Markos Zaharioudakis, Michael J. Carey

June 1997 **ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data SIGMOD '97**, Volume 26 Issue 2

Publisher: ACM Press

Full text available: pdf(1.81 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In this paper, we present four approaches to providing highly concurrent B+-tree indices in the context of a data-shipping, client-server OODBMS architecture. The first performs all index operations at the server, while the other approaches support varying degrees of client caching and usage of index pages. We have implemented the four approaches, as well as the 2PL approach, in the context of the SHORE OODB system at Wisconsin, and we present experimen ...

16 Database theory, technology and applications (DTTA): 1-2PC: the one-two phase atomic commit protocol



Yousef J. Al-Houmaily, Panos K. Chrysanthis

March 2004 **Proceedings of the 2004 ACM symposium on Applied computing SAC '04**

Publisher: ACM Press

Full text available: pdf(204.82 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper proposes a one-phase, two-phase commit (1-2PC) protocol that can be used to atomically commit Internet transactions distributed across sites in a wide area network. The 1-2PC protocol is characterized by its ability to dynamically select between one-phase and two-phase atomic commit protocols depending on the behavior of transactions and the system requirements. Thus, it offers the performance advantages of the one-phase atomic commit protocol whenever possible, while still providing ...

Keywords: Internet transactions, atomic commit protocol

17 ReVive: cost-effective architectural support for rollback recovery in shared-memory multiprocessors



Milos Prvulovic, Zheng Zhang, Josep Torrellas

May 2002 **ACM SIGARCH Computer Architecture News , Proceedings of the 29th annual international symposium on Computer architecture ISCA '02 , Proceedings of the 29th annual international symposium on Computer architecture ISCA '02**, Volume 30 Issue 2

Publisher: IEEE Computer Society, ACM Press

Full text available:

pdf(1.38 MB)

[Publisher Site](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper presents ReVive, a novel general-purpose rollback recovery mechanism for shared-memory multiprocessors. ReVive carefully balances the conflicting requirements of availability, performance, and hardware cost. ReVive performs checkpointing, logging, and distributed parity protection, all memory-based. It enables recovery from a wide class of errors, including the permanent loss of an entire node. To maintain high performance, ReVive includes specialized hardware that performs frequent o ...

Keywords: fault tolerance, shared-memory multiprocessors, rollback recovery, recovery, BER, logging, parity, checkpointing, availability

18 Crash recovery in client-server EXODUS

Michael J. Franklin, Michael J. Zwillig, C. K. Tan, Michael J. Carey, David J. DeWitt
June 1992 **ACM SIGMOD Record , Proceedings of the 1992 ACM SIGMOD international conference on Management of data SIGMOD '92**, Volume 21 Issue 2
Publisher: ACM Press

Full text available:  pdf(1.50 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In this paper, we address the correctness and performance issues that arise when implementing logging and crash recovery in a page-server environment. The issues result from two characteristics of page-server systems: 1) the fact that data is modified and cached in client database buffers that are not accessible by the server, and 2) the performance and cost trade-offs that are inherent in a client-server environment. We describe a recovery system that we have implemented for the client-ser ...

19 Paper session DB-1 (databases): networks and peer-to-peer: Decentralized coordination of transactional processes in peer-to-peer environments

Klaus Haller, Heiko Schuldt, Can Türker
October 2005 **Proceedings of the 14th ACM international conference on Information and knowledge management CIKM '05**
Publisher: ACM Press


Full text available:  pdf(369.47 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Business processes executing in peer-to-peer environments usually invoke Web services on different, independent peers. Although peer-to-peer environments inherently lack global control, some business processes nevertheless require global transactional guarantees, i.e., atomicity and isolation applied at the level of processes. This paper introduces a new decentralized serialization graph testing protocol to ensure concurrency control and recovery in peer-to-peer environments. The uniqueness of t ...

Keywords: DSGT, decentralized coordination, global correctness, partial rollback, peer-to-peer communication, transactional processes

20 A new approach to developing and implementing eager database replication protocols

Bettina Kemme, Gustavo Alonso
September 2000 **ACM Transactions on Database Systems (TODS)**, Volume 25 Issue 3
Publisher: ACM Press

Full text available:  pdf(449.43 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Database replication is traditionally seen as a way to increase the availability and performance of distributed databases. Although a large number of protocols providing data consistency and fault-tolerance have been proposed, few of these ideas have ever been used in commercial products due to their complexity and performance implications. Instead, current products allow inconsistencies and often resort to centralized approaches which eliminates some of the advantages of replication. As an ... nevertheless require global transacti

Keywords: database replication, fault-tolerance, group communication, isolation levels, one-copy-serializability, replica control, total error multicast

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.

Dial g DataStar

options

logoff

feedback

help

databases

search
page

titles

Document

Select the documents you wish to save or order by clicking the box next to the document, or click the link above the document to order directly.

save

locally as: PDF document



search strategy: do not include the search strategy



order

copy to
Clipboard
☒ Select All

- 1 Design of a MidO2PL database replication protocol in the MADIS middlewa
- 2 An integrated commit protocol for mobile network databases.
- 3 Correctness of parallel executions in multidatabase systems ruled by strict 2
- 4 Things every update replication customer should know.
- 5 Two-phase commit optimizations and tradeoffs in the commercial environ
- 6 On the heterogeneity of distributed databases integrating commit protocol
- 7 Commit protocols for externalized-commit heterogeneous database system
- 8 Transaction recovery in federated autonomous databases.
- 9 Open commit protocols tolerating commission failures.
- 10 Distributed scheduling of meetings: a case study in prototyping distribut
- 11 Multimedia information system platforms.
- 12 Efficient consistency maintenance in dispersive distributed databases.
- 13 A global checkpointing scheme for recovery in distributed database systems.

☒ document 1 of 13 Order Document

Inspec - 1898 to date (INZZ)

Accession number & update

0008995788 20070101.

Title

Design of a MidO2PL **database** replication protocol in the MADIS middleware architecture.

Conference information

Proceedings. 20th International Conference on Advanced Information Networking and Applications, Vienna, Austria, 18-20 April 2006.

Source

Proceedings. 20th International Conference on Advanced Information Networking and Applications, 2006, p. 5 pp., 10 refs, pp. CD-ROM, ISBN: 0-7695-2466-4.
Publisher: IEEE computer Society, Los Alamitos, CA, USA.

Author(s)

Armendariz-Inigo-J-E, Munoz-Escoi-F-D, Garitagoitia-J-R, de-Mendivil-J-R-G.

Author affiliation

Armendariz-Inigo, J.E., Univ. Publica de Navarra, Spain.

Abstract

Middleware **database** replication techniques is a way to increase performance and fault tolerance without modifying the **database** management system (DBMS) internals. However, it introduces an additional overhead that may lead to poor response times. In this paper we present a modification of the optimistic two **phase** locking (O2PL) M.J. Carey et al. (1991) protocol that orders transactions by way of a deadlock prevention schema, instead of using the total order transaction delivery obtained by

group communication systems (GCSs) G. Chockler et al. (2001) techniques, and do not need the **2 phase commit** (2PC) rule P.A. Bernstein et al. (1987). We formalize its definition as a state transition system and show that it is 1-copy-serializable (1CS) P.A. Bernstein et al. (1987).

Descriptors

FAULT-TOLERANCE; MIDDLEWARE; OPTIMISATION; PROTOCOLS; REPLICATED-DATABASES.

Classification codes

C6150N Distributed-systems-software*;
C6160B Distributed-databases;
C5640 Protocols;
C1180 Optimisation-techniques.

Keywords

optimistic-two-phase-locking; **MidO2PL-database-replication-protocol**; **MADIS-middleware-architecture**; **fault-tolerance**; **database-management- system**; **DBMS**; **deadlock-prevention-scheme**.

Treatment codes

P Practical;
T Theoretical-or-mathematical.

Language

English.

Publication type

Conference-paper.

Availability

CCCC: 1550-445X/2006/\$20.00.

Publication year

2006.

Publication date

20060000.

Edition

2006029.

Copyright statement

Copyright 2006 The Institution of Engineering and Technology.

(c) 2007 The Institution of Engineering and Technology

MidO2PL-database-replication-protocol; **MADIS-middleware-architecture**; **fault-tolerance**; **database-management- system**; **DBMS**; **deadlock-prevention-scheme**

☒ **document 2 of 13** Order Document

Inspec - 1898 to date (INZZ)**Accession number & update**

0008670377 20070101.

Title

An integrated **commit** protocol for mobile network databases.

Conference information

Proceedings. 9th International **Database** Engineering and Applications Symposium (IDEAS 2005), Montreal, Que., Canada, 25-27 July 2005.

Source

Proceedings. 9th International **Database** Engineering and Applications Symposium (IDEAS 2005), 2005, p. 244-50, 12 refs, pp. xiii+456, ISBN: 0-7695-2404-4.
 Publisher: IEEE Comput. Soc, Los Alamitos, CA, USA.

Author(s)

Bose-J-H, Obermeier-S, Bottcher-S, Schwepe-H, Gruenwald-L, Steenweg-T.
 Editor(s): Desai-B-C, Vossen-G.

Author affiliation

Bose, J.-H., Inst. of Comput. Sci., Freie Univ. Berlin.

Abstract

While traditional fixed-wired network protocols like **2-phase-commit** guarantee atomicity, we cannot use them in mobile low bandwidth networks where network partitioning, node failure, and message

loss may result in blocking. To deploy traditional **database** applications easily into a mobile environment, there is a demand for a protocol which guarantees an atomic **commit** of transactions. This paper introduces a protocol which can guarantee such atomic commitment in mobile environments using a combination of **commit** and consensus protocols. In addition, it takes advantage of mobile network sub-structures like single-hop environments to reduce message transfer costs.

Descriptors

 DATABASE-MANAGEMENT-SYSTEMS;  MOBILE-COMPUTING;  PROTOCOLS.

Classification codes

B6250F Mobile-radio-systems*;
B6210L Computer-communications;
B6150M Protocols;
C5620 Computer-networks-and-techniques*;
C6160 Database-management-systems-DBMS;
C5640 Protocols.

Keywords

integrated-commit-protocol; mobile-network-databases; mobile-low- bandwidth-networks;
network-partitioning; consensus-protocol; message-transfer-cost-reduction.

Treatment codes

P Practical.

Language

English.

Publication type

Conference-paper.

Availability

CCCC: 0 7695 2404 4/2005/\$20.00.

Publication year

2005.

Publication date

20050000.

Edition

2005048.

Copyright statement

Copyright 2005 IEE.

(c) 2007 The Institution of Engineering and Technology

☒ **document 3 of 13** Order Document

Inspec - 1898 to date (INZZ)

Accession number & update

0005478976 20070101.

Title

Correctness of parallel executions in multidatabase systems ruled by strict **2 phase** locking.

Conference information

Proceedings of 9th International Conference on Parallel and Distributed Computing Systems. PDCS '96, Dijon, France, 25-27 Sept. 1996.

Sponsor(s): ISCA; IEEE Comput. Soc; IEEE Tech. Committee on Operating Syst; et al.

Source

Proceedings of the ISCA International Conference. Parallel and Distributed Computing Systems, 1996, **vol.2**, p. 618-25 **vol.2**, 6 refs, pp. 2 vol. x+825.

Publisher: Int. Soc. Comput. & Their Appl.-ISCA, Raleigh, NC, USA.

Author(s)

Besancenot-J, Cart-M, Ferrie-J, Morpain-C, Pons-J-F, Pucheral-P.

Editor(s): Yetongnon-K, Hariri-S.








Author affiliation

Besancenot, J., PRISM Lab., Versailles Univ., France.

Abstract

This paper addresses the problem of transaction management in multidatabase systems where the participating local DBMSs serialize transactions by **2 phase-locking (2PL)** and synchronize their **commit** actions through a **2 phase-commit (2PC)** protocol. These DBMSs represent the majority of commercial relational and object-oriented DBMSs. We demonstrate that if local DBMSs support strict 2PL optimization (i.e., relax the read locks during the prepare **phase** of the 2PC protocol), then schedules of parallel global transactions may become non-globally serializable. X/Open distributed transaction processing protocol avoids the problem by adding a blocking synchronization **phase** at transaction end, thereby losing the whole benefit of strict 2PL. This paper proposes two strategies to preserve the benefit of strict 2PL while ensuring correct executions.

Descriptors

 DATABASE-THEORY;  OBJECT-ORIENTED-DATABASES;  PARALLEL-PROCESSING;
 PROTOCOLS;  RELATIONAL-DATABASES;  SYNCHRONISATION;  VERY-LARGE-DATABASES.

Classification codes

C4250 Database-theory*;
C6160D Relational-databases;
C6160J Object-oriented-databases;
C6150N Distributed-systems-software;
C5640 Protocols.

Keywords

parallel-execution-correctness; multidatabase-systems; **two-phase- locking**; distributed-transaction-processing-protocol; synchronization; **relational-database**; **object-oriented-database**; optimization; X/Open- protocol; transaction-management.

Treatment codes

P Practical;
I Theoretical-or-mathematical.

Language

English.

Publication type

Conference-paper.

Publication year

1996.

Publication date

19960000.

Edition

1997003.

Copyright statement

Copyright 1997 IEE.

(c) 2007 The Institution of Engineering and Technology

Full text available at



USPTO Full Text Retrieval Options

☒ **document 4 of 13** Order Document

Inspec - 1898 to date (INZZ)

Accession number & update

0004999598 20070101.

Title

Things every update replication customer should know.

Conference information

1995 ACM SIGMOD International Conference on Management of Data, San Jose, CA, USA, 22-25 May 1995.

Sponsor(s): ACM.

Source

SIGMOD Record, {SIGMOD-Rec-USA}, June 1995, vol. 24, no. 2, p. 439-40, 4 refs, CODEN: SRECD8, ISSN: 0163-5808, USA.

Author(s)

Goldring-R.

Author affiliation

Goldring, R., Santa Teresa Lab., IBM Corp., San Jose, CA.

Abstract

Replication in distributed **database** systems is discussed. Replication makes transparent the problems of remote access delays and the management of data redundancy. The commercial market for distributed **database** features has been slowly building over the years, beginning with simple remote access gateways. Today, replication appears to deliver on the 1980s ideal, with a robust asynchronous infrastructure. Current commercial technology though, continues to fall short of that ideal.

"Asynchronous replication" is a pleasant term to describe the operation of a distributed **database** running without concurrency control. In practice, DBMSs which use locking mechanisms in local operation are connected into replication networks without benefit of a global serialization mechanism, such as a synchronous **2-phase commit** protocol. The notion of a transaction is thus compromised. Four properties, atomicity, consistency, isolation and durability (ACID), have come to define a transaction system. With asynchronous replication, there is no isolation of transactions. Transactions run in parallel without any guarantee that a transaction sees the most current state of the **database** before making an update. Updates then, are not serialized. One of the many benefits derived from the ACID properties is a serial history of transaction execution, an absolute necessity to satisfy audit requirements in regulated industries. Without a serial history, it is impossible to reliably state who updated a **database** from state N to state N+1. Not all replication systems guarantee a serial history.

Descriptors

CONCURRENCY-CONTROL; DISTRIBUTED-DATABASES; REDUNDANCY; SYSTEM-RECOVERY; TRANSACTION-PROCESSING.

Classification codes

C6160B Distributed-databases*;
C6150G Diagnostic-testing-debugging-and-evaluating-systems;
C6130 Data-handling-techniques;
C4250 Database-theory;
C6150N Distributed-systems-software.

Keywords

update-replication-customer; **distributed-database-systems**; remote-access-delays; data-redundancy; commercial-market; remote-access-gateways; robust-asynchronous-infrastructure; asynchronous-replication; locking-mechanisms; local-operation; global-serialization-mechanism; ACID-properties; atomicity; consistency; isolation; durability; transaction-system; transaction-execution; serial-history.

Treatment codes

P Practical.

Language

English.

Publication type

Conference-paper; Journal-paper.

Publication year

1995.

Publication date

19950600.

Edition

1995028.

Copyright statement

Copyright 1995 IEE.

(c) 2007 The Institution of Engineering and Technology

Inspec - 1898 to date (INZZ)**Accession number & update**

0004810770 20070101.

Title**Two-phase commit** optimizations and tradeoffs in the commercial environment.**Conference information**

Proceedings of IEEE 9th International Conference on Data Engineering, Vienna, Austria, 19-23 April 1993.

Sponsor(s): IEEE Comput. Soc. Tech. Committee on Data Eng; Austrian Comput. Soc. (OCG); Gi; ERCIM; INRIA; GMD; CWI; RAL FAW-A.

Source

Proceedings. Ninth International Conference on Data Engineering (Cat. No.92CH3258-1), 1993, p. 520-9, 31 refs, pp. xviii+687, ISBN: 0-8186-3570-3.

Publisher: IEEE Comput. Soc. Press, Los Alamitos, CA, USA.

Author(s)Samaras-G, Britton-K, Citron-A, Mohan-C.**Author affiliation**

Samaras, G., Britton, K., Citron, A., Distributed Syst. Archit. IBM Almaden Res. Center, San Jose, CA, USA.

AbstractEleven **two-phase commit** (2PC) protocol variations that optimize towards the normal case are described and compared with a baseline 2PC protocol. Environments in which they are most effective are discussed. The variations are compared and contrasted in terms of number of message flows, number of log writes (both forced and non-forced), probability of heuristic damage, how damage is reported, and other tradeoffs.**Descriptors** DATA-INTEGRITY;  OPTIMISATION;  PROTOCOLS;  TRANSACTION-PROCESSING.**Classification codes**C6160 Database-management-systems-DBMS*.C5640 Protocols.**Keywords****2-phase-commit-protocols**; damage-reporting; forced-writes; unforced-writes; optimizations; tradeoffs; commercial-environment; message-flows; log-writes; heuristic-damage.**Treatment codes**P Practical.**Language**

English.

Publication typeConference-paper.**Availability**

CCCC: 1063-6382/93/\$03.00.

Digital object identifier

10.1109/ICDE.1993.344028.

Publication year

1993.

Publication date

19930000.

Edition

1994045.

Copyright statement

Copyright 1994 IEE.

(c) 2007 The Institution of Engineering and Technology

☒ **document 6 of 13** Order Document

Inspec - 1898 to date (INZZ)**Accession number & update**

0004731632 20070101.

TitleOn the heterogeneity of distributed databases integrating **commit** protocols.**Conference information**

14th International Conference on Distributed Computing Systems, Pozman, Poland, 21-24 June 1994.

Sponsor(s): IEEE Comput. Soc. Tech. Committee on Distributed Process; Polish Inf. Process. Soc.

Source

Proceedings of the 14th International Conference on Distributed Computing Systems (Cat. No.94CH3450-4), 1994, p. 380-6, 16 refs, pp. xix+651, ISBN: 0-8186-5840-1.

Publisher: IEEE Comput. Soc. Press, Los Alamitos, CA, USA.

Author(s)Papadopoulos-C.V.**Author affiliation**

Papadopoulos, C.V., Dept. of Comput. Sci., Piraeus Univ., Greece.

Abstract

Explores atomic commitment among heterogeneous distributed databases. Two cases are considered, one where the individual databases participating in the system do not externalize their **commit** protocols, and one where the participating databases do externalize their **commit** protocols. With regard to the latter case, I propose new heterogeneous **commit** protocols, one for joining databases that have different **2-phase commit** protocols, and one for uniting databases which employ a **3-phase** protocol with others employing a **2-phase commit** protocol.

DescriptorsDISTRIBUTED-DATABASES; PROTOCOLS; TRANSACTION-PROCESSING.**Classification codes**C6160B Distributed-databases*;C5640 Protocols.**Keywords**

heterogeneous-distributed-databases; atomic-commitment; **commit**- protocol-externalization; **heterogeneous-commit-protocols**; **database-joining**; **database-uniting**; **2-phase-protocol**; **3-phase-protocol**.

Treatment codesP Practical.**Language**

English.

Publication typeConference-paper.**Availability**

CCCC: 1063-6927/94/03.00.

Digital object identifier

10.1109/ICDCS.1994.302440.

Publication year

1994.

Publication date

19940000.

Edition

1994030.

Copyright statement

Copyright 1994 IEE.

(c) 2007 The Institution of Engineering and Technology

Full text available at **USPTO Full Text Retrieval Options**☒ **document 7 of 13** Order Document

Inspec - 1898 to date (INZZ)**Accession number & update**

0004673032 20070101.

Title**Commit protocols for externalized-commit heterogeneous database systems.****Source**

Distributed and Parallel Databases, {Distrib-Parallel-Databases-Netherlands}, April 1994, vol. 2, no. 2, p. 209-34, 20 refs, CODEN: DPADEH, ISSN: 0926-8782, Netherlands.

Author(s)

Tal-A, Alonso-R.

Author affiliation

Tal, A., Dept. of Comput. Sci., Princeton Univ., NJ.

Abstract

In designing a heterogeneous **database** systems, one of the main technical challenges is developing techniques for ensuring global **commit**. That is, guaranteeing that a transaction spanning multiple individual **database** management systems (DBMSs) either commits at all the participating DBMSs or at none of them. Previous work in this area typically assumes that the participating DBMSs do not provide a mechanism for interacting with their **commit** facilities. While this is true in many cases, in practice there are systems which support a programmatic interface to their **commit** protocols. We refer to **database** systems offering such facilities as externalized **commit** DBMSs. The focus of this paper is on **commit** protocols for these systems. We propose two new **commit** protocols for externalized **commit** DBMSs. The first may be used to obtain global **commit** in heterogeneous **database** systems composed of DBMSs with different **2-phase commit** protocols (e.g., centralized and linear). The second protocol is more general, and ensures global **commit** even if the participating DBMSs employ **3-phase commit** protocols. The more general protocol also preserves **database** autonomy, since it does not block a DBMS upon failure of another system. We describe both protocols in detail and prove their correctness.

Descriptors**CONCURRENCY-CONTROL; DISTRIBUTED-DATABASES; TRANSACTION-PROCESSING.****Classification codes**

C6160B Distributed-databases*;

C6150N Distributed-systems-software.

Keywords**commit-protocols; externalized-commit-heterogeneous-database-systems; global-commit; database-management-systems; 2-phase-commit; database- autonomy; correctness.****Treatment codes**

P Practical.

Language

English.

Publication type

Journal-paper.

Availability

CCCC: 0926-8782/94/\$5.00.

Publication year

1994.

Publication date

19940400.

Edition

1994018.

Copyright statement

Copyright 1994 IEE.

(c) 2007 The Institution of Engineering and Technology.

Full text available at **USPTO Full Text Retrieval Options**

☒ **document 8 of 13** Order Document

Inspec - 1898 to date (INZZ)

Accession number & update

0004673030 20070101.

Title

Transaction recovery in federated autonomous databases.

Source

Distributed and Parallel Databases, {Distrib-Parallel-Databases-Netherlands}, April 1994, vol. 2, no. 2, p. 151-82, 26 refs, CODEN: DPADEH, ISSN: 0926-8782, Netherlands.

Author(s)

San-Yih-Hwang, Srivastava-J, Jianzhong-Li.

Author affiliation

San-Yih Hwang, Srivastava, J., Jianzhong Li, Dept. of Comput. Sci., Minnesota Univ., Minneapolis, MN, USA.

Abstract

Transaction management for federated **database** systems (FDBSs), where each participating DBMS is autonomous, supports global transactions that can access more than one **database**. A number of concurrency control algorithms exist for the failure-free environment. Handling transaction failure, due to concurrency control or site related reasons, becomes especially difficult in the presence of autonomy. Due to local autonomy, carrying out **2-phase commit** may not be possible. This can be simulated by providing a server on top of the pre-existing DBMS at each site, which is responsible for submitting the local operations to the associated DBMS and communicating with the transaction's originating site. In this paper we formalize the problem of ensuring transaction consistency in an FDBS environment in the presence of failure. The key problem is that due to autonomy, the local DBMS and FDBS may have different views of an execution sequence generated at a site. Local recoverability is identified as the property of local execution sequences necessary for correctness. The other main problem is of guaranteeing that the various locally recoverable histories are consistent with each other. These identified properties are necessary and sufficient conditions for the correctness of FDBS recovery algorithms, and can be used to evaluate the correctness of the proposed algorithms. This paper also presents an FDBS recovery algorithm that has been proved to be correct. Formal proofs of all properties and a comparison of different algorithms are provided.

Descriptors

☒ CONCURRENT-CONTROL; ☒ DISTRIBUTED-DATABASES; ☒ SYSTEM-RECOVERY;
☒ TRANSACTION-PROCESSING.

Classification codes

C6160B Distributed-databases*;
C6150N Distributed-systems-software.

Keywords

transaction-recovery; federated-autonomous-databases; federated- **database-systems**; DBMS;
global-transactions; concurrency-control- algorithms; **2-phase-commit**; transaction-consistency;
local-recoverability; local-execution-sequences; correctness; recoverable-histories.

Treatment codes

P Practical;
I Theoretical-or-mathematical.

Language

English.

Publication type

Journal-paper.

Availability

CCCC: 0926-8782/94/\$5.00.

Publication year

1994.

Publication date

19940400.

Edition

1994018.

Copyright statement

Copyright 1994 IEE.

(c) 2007 The Institution of Engineering and Technology

Full text available at



USPTO Full Text Retrieval Options

☒ document 9 of 13 [Order Document](#)

Inspec - 1898 to date (INZZ)

Accession number & update

0004503935 20070101.

Title

Open **commit** protocols tolerating commission failures.

Source

ACM Transactions on **Database Systems**, {**ACM-Trans-Database-Syst-USA**}, June 1993, vol. 18, no. 2, p. 289-332, 31 refs, CODEN: ATDSD3, ISSN: 0362-5915, USA:

Author(s)

[Rothermel-K](#), [Pappe-S](#).

Author affiliation

Rothermel, K., Inst. of Parallel & Distributed High Performance Syst., Stuttgart Univ., Germany.

Abstract

To ensure atomicity of transactions in distributed systems so-called **2-phase commit** (2PC) protocols have been proposed. This paper proposes a family of 2PC protocols that tolerate any number of omission failures at trusted nodes and any number of commission and omission failures at nontrusted nodes. The proposed protocols ensure that (at least) the trusted nodes participating in a transaction eventually terminate that transaction in a consistent manner. Unlike Byzantine **commit** protocols, the protocols do not incorporate mechanisms for achieving Byzantine agreement, which has advantages in terms of complexity: the protocols have the same or only a slightly high message complexity than traditional 2PC protocols.

Descriptors

[DISTRIBUTED-DATABASES](#); [FAULT-TOLERANT-COMPUTING](#); [PROTOCOLS](#);
 [TRANSACTION-PROCESSING](#).

Classification codes

[C6160B Distributed-databases*](#);
[C5640 Protocols](#);
[C6155 Computer-communications-software](#).

Keywords

two-phase-commit-protocols; transaction-atomicity; fault-tolerant- computing; commission-failures; distributed-systems; 2PC-protocols; omission-failures; **Byzantine-commit-protocols**; message-complexity.

Treatment codes

[P Practical](#).

Language

English.

Publication type

[Journal-paper](#).

Availability

CCCC: 0362-5915/93/0600-0316\$01.50.

Publication year

1993.

Publication date

19930600.

Edition

1993041.

Copyright statement

Copyright 1993 IEE.

(c) 2007 The Institution of Engineering and Technology

☒ **document 10 of 13** Order Document

Inspec - 1898 to date (INZZ)

Accession number & update

0004387205 20070101.

Title

Distributed scheduling of meetings: a case study in prototyping distributed applications.

Conference information

ICSI '92. Proceedings of the Second International Conference on Systems Integration (Cat. No.92TH0444-0), Morristown, NJ, USA, 15-18 June 1992.

Sponsor(s): IEEE; New Jersey Inst. Technol; ACM.

Source

ICSI '92. Proceedings of the Second International Conference on Systems Integration (Cat. No.92TH0444-0), 1992, p. 656-65, 15 refs, pp. xx+746, ISBN: 0-8186-2697-6.

Publisher: IEEE Comput. Soc. Press, Los Alamitos, CA, USA.

Author(s)

Biswas-J, Bhonsle-S, Tan-Chee-Wee, Tay-Sen-Yong, Wang-Weiguo.

Editor(s): Ng-P-A, Seifert-L-C, Ramamoorthy-C-V, Yeh-R-T.

Author affiliation

Biswas, J., Bhonsle, S., Tan Chee Wee, Tay Sen/Yong, Wang Weiguo, Inst. of Syst. Sci., Nat. Univ. of Singapore, Singapore.

Abstract

The authors have developed a combined meeting-scheduling cum calendar-management system called CAMEL, that eliminates the tedium of scheduling a meeting. CAMEL uses a hunting feature that enables the tracking of user logins. The approach towards software development is generative as well as declarative, through extensive use of toolkits and reusable software. The authors describe the essential ingredients of CAMEL. It is a fairly complex distributed application using a distributed **database** and distributed user related information such as preference parameters. Tools from RAPIDS toolkit, especially remote procedure call subsystem, n-party interaction subsystem, are heavily used to produce this application. It also uses many services provided by RAPIDS, such as name server, user information server etc. To maintain consistency of distributed data the application makes use of the **2-phase commit** and **2-phase locking** primitives provided by RAPIDS.

Descriptors

DISTRIBUTED-DATABASES; DISTRIBUTED-PROCESSING; PERSONAL-COMPUTING; SCHEDULING; SOFTWARE-PROTOTYPING; SOFTWARE-TOOLS.

Classification codes

C7100 Business-and-administration*;

C6110 Systems-analysis-and-programming;

C6115 Programming-support;

C6150N Distributed-systems-software;

C6160B Distributed-databases.

Keywords

user-logic-tracking; distributed-scheduling; software-prototyping; distributed-applications; meeting-scheduling; calendar-management-system; CAMEL; hunting-feature; software-development; reusable-software; **distributed-database**; distributed-user-related-information; preference-parameters; RAPIDS-toolkit; remote-procedure-call-subsystem; n-party-interaction-subsystem; name-server; user-information-server; **2-phase-commit**; **2-phase-locking-primitives**.

Treatment codes

P Practical.

Language

English.

Publication type

Conference-paper.

Availability

CCCC: 0 8186 2697 6/92\$03.00.

Digital object identifier

10.1109/ICSI.1992.217265.

Publication year

1992.

Publication date

19920000.

Edition

1993014.

Copyright statement

Copyright 1993 IEE.

(c) 2007 The Institution of Engineering and Technology

Full text available at **USPTO Full Text Retrieval Options**☒ **document 11 of 13** [Order Document](#)**Inspec - 1898 to date (INZZ)****Accession number & update**

0004345892 20070101..

Title

Multimedia information system platforms.




Source

Korea Information Science Society Review, {Korea-Inf-Sci-Soc-Rev-South-Korea}, 1992, vol. 10, no. 5, p. 5-9, 26 refs, CODEN: CHKWEN, ISSN: 1015-9908, South Korea.

Author(s)Kyu-Yong-Whang.**Abstract**

Discusses multimedia information retrieval systems and document systems; hypertext and hypermedia; office information systems; engineering, scientific and geographical databases; remote information access; mixed mail services; electronic publishing; KAIST; VTRs and laserdisc players; Hyper ODA; SGML; authorization; data manipulation languages; **2-phase commit**; long data items; binary large data objects; the IBM Audio-Visual Connection system; Multimedia Toolbox; DVI; Authology; MacroMind; multimedia authoring; reference architecture; persistence; navigation; and the work of the Multimedia and Hypermedia Information Coding Expert Group (MHEG).

Descriptors

 DATABASE-MANAGEMENT-SYSTEMS;  INFORMATION-RETRIEVAL-SYSTEMS;
 MULTIMEDIA-SYSTEMS.

Classification codes

C6160S Spatial-and-pictorial-databases*;
C7250 Information-storage-and-retrieval.

Keywords

multimedia-information-system-platforms; information-retrieval-systems; document-systems; hypertext; hypermedia; office-information-systems; databases; remote-information-access; mail-services; electronic-publishing; KAIST; Hyper-ODA; SGML; authorization; data-manipulation-languages; multimedia-authoring; reference-architecture; Multimedia-and-Hypermedia-Information-Coding-Expert-Group.

Treatment codesG General-or-review.**Language**

Korean.

Publication typeJournal-paper.**Publication year**

1992.

Publication date

19920000.

Edition

1993006.

Copyright statement

Copyright 1993 IEE.

(c) 2007 The Institution of Engineering and Technology

Full text available at **USPTO Full Text Retrieval Options**☒ **document 12 of 13** [Order Document](#)**Inspec - 1998 to date (INZZ)****Accession number & update**

0004171962 20070101.

Title

Efficient consistency maintenance in dispersive distributed databases.

Source

Informationstechnik - IT, {Inftech-IT-Germany}, April 1992, vol. 34, no. 2, p. 132-9, 25 refs, CODEN: ICSAES, ISSN: 0179-9738, Germany.

Author(s)[Rautenstrauch-C.](#)**Author affiliation**

Rautenstrauch, C., Westfälische Wilhelms-Univ., Munster, Germany.

Abstract

Today almost every **database** system vendor claims data distribution as a strategic goal. The technical base for distributed data management consists of the **2-phase-commit** protocol (2PC). Unfortunately the 2PC contains weaknesses in terms of security and performance. The paper presents the application programmers point of view how to maintain consistency for multi-user environments without applying 2PC. At first several topics of distributed **database** environments are reviewed, further on a method for consistency maintenance (called update-by-need) is introduced, analyzed and illustrated by an example.

Descriptors [DISTRIBUTED-DATABASES.](#)**Classification codes**[C6160B Distributed-databases*.](#)**Keywords**consistency-maintenance; dispersive-distributed-databases; data-distribution; distributed-data-management; **2-phase-commit-protocol**; 2PC; security; performance; multi-user-environments; distributed- **database-environments**; update-by-need.**Treatment codes**[P Practical.](#)**Language**

German.

Publication type[Journal-paper.](#)**Publication year**

1992.

Publication date

19920400.

Edition

1992028.

Copyright statement

Copyright 1992 IEE.

(c) 2007 The Institution of Engineering and Technology

Full text available at **USPTO Full Text Retrieval Options**

☒ **document 13 of 13** Order Document

Inspec - 1898 to date (INZZ)

Accession number & update

0003400077 20070101.

Title

A global checkpointing scheme for recovery in distributed **database** systems.

Source

Journal of the Korea Information Science Society, {J-Korea-Inf-Sci-Soc-South-Korea}, Jan. 1989, vol. 16, no. 1, p. 77-83, 20 refs, CODEN: HJKHDC, ISSN: 0258-9125, South Korea.

Author(s)

Ju-Gyun-Kim, Kern-Koh.

Abstract

The demand for highly reliable computer systems has led to techniques for construction of a consistent state after transient failures. In DDBMS 'checkpointing and rollbackrecovery' are well known techniques that allow processes to make progress in spite of failures. The paper presents and discusses a periodic checkpointing method coordinated with the **2-phase commit** protocol which ensures global consistency and recoverability when data are fully replicated. In contrast to previous methods, this method not only solves domino-effect problems but also needs no additional hardware and communication overhead. It can be easily implemented in any systems using the **2-phase commit** protocol.

Descriptors

DISTRIBUTED-DATABASES.

Classification codes

C6160B Distributed-databases*.

Keywords

global-checkpointing-scheme; recovery; **distributed-database-systems**; reliable-computer-systems; consistent-state; transient-failures; DDBMS; rollbackrecovery; periodic-checkpointing-method; **2-phase-commit-protocol**; global-consistency; recoverability; domino-effect- problems.

Treatment codes

P Practical.

Language

Korean.

Publication type

Journal-paper.

Publication year

1989.

Publication date

19890100.

Edition

1989015.

Copyright statement

Copyright 1989 IEE.

(c) 2007 The Institution of Engineering and Technology

save

locally as: **PDF document**

search strategy: **do not include the search strategy**

order

copy to
Clipboard

distributed-database-systems; reliable

Top - News & FAQs - Dialog

ProQuest

[Return to the USPTO NPL Page](#) | [Help](#)

Basic

Advanced

Topics

Publications

My Research
0 marked items

Interface language:

English



































Databases selected: Multiple databases...

Results

72 documents found for: 2 phase commit and database >> [Refine Search](#) | [Set Up Alert](#) [All sources](#) | [Scholarly Journals](#) | [Trade Publications](#) | [Newspapers](#) | [Dissertations](#)☐ Mark all 0 marked items: [Email](#) / [Cite](#) / [Export](#) [Show only full text](#) Sort results by: [Most relevant](#)

- ☐ 1. **Consensus on transaction commit**
Jim Gray, Leslie Lamport. *ACM Transactions on Database Systems*. New York: Mar 2006. Vol. 31, Iss. 1
 [Link to full text](#) [Abstract](#)
- ☐ 2. **Transaction processing in mobile database system**
by Prabhu, Nitin, Ph.D., University of Missouri - Kansas City, 2006, 157 pages; AAT 3245280
 [Abstract](#) [24 Page Preview](#) [Full Text - PDF](#) [Order a](#)
- ☐ 3. **NEON SYSTEMS: ShadowConnect offers high performance mainframe integration at connector-based Low cost entry point for discrete distributed computing projects; Priced to fit within the ROI budget application development projects; J2EE connectors are certified for BEA Weblogic**
M2 Presswire. Coventry: Nov 20, 2002. p. 1
 [Full text](#) [Abstract](#)
- ☐ 4. **NEON SUPPORTS IBM WEBSHERE WITH JDBC ACCESS TO OS/390**
Mainframe Computing. Boynton Beach: Jan 1, 2001. Vol. 14, Iss. 1; p. 1
 [Full text](#) [Abstract](#)
- ☐ 5. **Transaction management under J2EE 1.2; Explore your options for managing transactions**
Sanjay Mahapatra. *JavaWorld*. San Francisco: Jul 20, 2000. p. 1
 [Full text](#) [Abstract](#)
- ☐ 6. **New standard benefits**
Paul Rampel. *Enterprise Systems Journal*. Dallas: Sep 1999. Vol. 14, Iss. 9; p. 74 (1 page)
 [Abstract](#) [Abstract](#)
- ☐ 7. **StarQuest Software Adds Significant New Features to StarSQL; New Two-Phase Commit Capability Data Integrity, UNIX Support Added**
Business Editors/High-Tech Writers. *Business Wire*. New York: Jun 8, 1999. p. 1
 [Full text](#) [Abstract](#)
- ☐ 8. **New middleware brings PCs and host databases together**
John Cox. *Network World*. Framingham: May 3, 1999. Vol. 16, Iss. 18; p. 34 (1 page)
 [Full text](#) [Full Text - PDF](#) [Abstract](#)
- ☐ 9. **Achieving a multidatabase object transaction service through an asynchronous socket**
by McLaughlin, Michael James, D.C.S., Colorado Technical University, 1999, 203 pages; AAT 9953695
 [Abstract](#) [24 Page Preview](#) [Full Text - PDF](#) [Order a](#)
10. **ObjectStore 5.1 stands out for data management**

10/726,700

- ☐ **Maggie Biggs. InfoWorld.** San Mateo: Jul 13, 1998. Vol. 20, Iss. 28; p. 107 (1 page)
 [Text+Graphics](#)  [Full Text - PDF](#)  [Abstract](#)
- ☐ 11. **Netiva misses the mark**
Tom Stearns. Network World. Framingham: Dec 15, 1997. Vol. 14, Iss. 50; p. 116 (1 page)
 [Text+Graphics](#)  [Full Text - PDF](#)  [Abstract](#)
- ☐ 12. **Competitive Advantage: Comparing SQL servers... part 2**
BusinessWorld. Manila: Nov 25, 1997. p. NOPGCIT
 [Full text](#)  [Abstract](#)
- ☐ 13. **INFOTECH BUYLINE; [2 Edition]**
Dominion. Wellington, New Zealand: Apr 7, 1997. p. NOPGCIT
 [Full text](#)  [Abstract](#)
- ☐ 14. **Commit processing in distributed database systems and in heterogeneous multidatabase systems**
 by **Al-Houmailly, Yousef Jassem, Ph.D., University of Pittsburgh,** 1997, 172 pages; AAT 9736438
 [Abstract](#)  [24 Page Preview](#)  [Full Text - PDF](#)  [Order a](#)
- ☐ 15. **Design and evaluation of a consulting system for database design**
 by **Antony, Solomon Raj, Ph.D., Florida International University;** 1997, 175 pages; AAT 9805206
 [Abstract](#)  [24 Page Preview](#)  [Full Text - PDF](#)  [Order a](#)
- ☐ 16. **Two-phase commit: What it is and how it works**
Anonymous. Datamation. Barrington: Sep 1996. Vol. 42, Iss. 15; p. 64 (1 page)
 [Text+Graphics](#)  [Full Text - PDF](#)  [Abstract](#)
- ☐ 17. **Performance of a two-phase commit protocol**
Desai, Bipin C, Boutros, Boutros S. Information and Software Technology. Amsterdam: Sep 1996. Vol. 38, Iss. 9; p. 581 (19 pages)
 [Link to full text](#)  [Abstract](#)
- ☐ 18. **Oracle Webizes databases**
Anonymous. Network World. Framingham: Aug 12, 1996. Vol. 13, Iss. 33; p. 17 (1 page)
 [Full text](#)  [Full Text - PDF](#)  [Abstract](#)
- ☐ 19. **A nonblocking timed atomic commit protocol for distributed real-time database systems**
Wong, Johnny S K, Mitra, Soma. The Journal of Systems and Software. New York: Aug 1996. Vol. 34, Iss. 2; p. 103 (8 pages)
 [Link to full text](#)  [Abstract](#)
- ☐ 20. **Distributed operating systems**
Mullender, Sape J. ACM Computing Surveys. Baltimore: Mar 1996. Vol. 28, Iss. 1; p. 225 (3 pages)
 [Link to full text](#)  [Abstract](#)
- ☐ 21. **DITSE: An experimental distributed database system**
Ni, Cheng-Yao, Yuan, Shyan-Ming. Information and Software Technology. Amsterdam: Feb 1996. Vol. 38, Iss. 2; p. 103 (8 pages)
 [Link to full text](#)  [Abstract](#)
- ☐ 22. **An architecture and transaction model for large federated databases**
 by **Chazin, David, Ph.D., University of Illinois at Urbana-Champaign,** 1996, 192 pages; AAT 9712221

[Abstract](#) [24 Page Preview](#) [Full Text - PDF](#) [Order a](#)

- ☐ 23. **Transaktionsverwaltung in heterogenen, foederierten Datenbanksystemen**
by *Schaad, Werner*, Dr.sc.tech, Eidgenoessische Technische Hochschule Zuerich (Switzerland), 1996
AAT C538216

[Abstract](#)

- ☐ 24. **Commitment that counts**
Johnson, Johna Till. *Data Communications*. New York: Aug 1995. Vol. 24, Iss. 10; p. 54 (1 page)

[Full Text - PDF](#) [Abstract](#)

- ☐ 25. **Data replication: A risky Rx**
Ricciuti, Mike. *InfoWorld*. San Mateo: Jul 3, 1995. Vol. 17, Iss. 27; p. 1 (2 pages)

[Full text](#) [Full Text - PDF](#) [Abstract](#)

- ☐ 26. **Versant bolsters reliability of its object database**
Cole, Barb. *Network World*. Framingham: Jun 12, 1995. Vol. 12, Iss. 24; p. 35 (2 pages)

[Full text](#) [Full Text - PDF](#) [Abstract](#)

- ☐ 27. **Btrieve adds replication to Scalable SQL**
Ricciuti, Mike. *InfoWorld*. San Mateo: Apr 24, 1995. Vol. 17, Iss. 17; p. 25 (1 page)

[Full text](#) [Full Text - PDF](#) [Abstract](#)

- ☐ 28. **Personal Oracle and Oracle7 get replication**
Ricciuti, Mike. *InfoWorld*. San Mateo: Apr 17, 1995. Vol. 17, Iss. 16; p. 16 (1 page)

[Full text](#) [Full Text - PDF](#) [Abstract](#)

- ☐ 29. **QS/400 V3R1**
Dickey, Sam. *Midrange Systems*. Spring House: Mar 10, 1995. Vol. 8, Iss. 5; p. 20 (3 pages)

[Abstract](#)

- ☐ 30. **Partial information based integrity constraint checking**
by *Gupta, Ashish*, Ph.D., Stanford University, 1995, 162 pages; AAT 9516833

[Abstract](#) [24 Page Preview](#) [Full Text - PDF](#) [Order a](#)

1-30 of 72



< First | < Previous

Want to be notified of new results for this search? [Set Up Alert](#)

Results per

Basic Search

[Tools:](#) [Search Tips](#) [Browse Topics](#) [1 Recent Searches](#) [Search](#)[Clear](#)Database: [Select multiple databases](#)Date range: Limit results to: ☐ Full text documents only ☐ Scholarly journals, including peer-reviewed [About](#)

 [More Search Options](#) 

Copyright © 2007 ProQuest Information and Learning Company. All rights reserved. [Terms and Conditions](#)

[Text-only interface](#)



Copyright © 2007 ProQuest Information and Learning Company. All rights reserved. [Terms and Conditions](#)

[Text-only interface](#)